MHz. The base frequency (f_{base})for the network is defined as the lowest frequency of the Lower RF Band.

The Lower and Upper RF Bands are further subdivided into sub-bands. The first and last 0.5 MHz of each band are designated as guard bands and are hence unused. The remaining 4 MHz in each band is then subdivided into 4 Sub-bands sequentially numbered from 0 to 3. Each Sub-band contains a set of frequencies in the Lower RF Band and another set of frequencies in the Upper RF Band. The extension L indicates the set within the Lower RF Band and U indicates the set within the Upper RF Band.

In one embodiment there are a total of 2560 frequency tones equally spaced in the 8 MHz of available bandwidth. There are 1280 tones in each Band, and 640 tones in each Sub-band (320 frequencies in the lower band and 320 frequencies in the upper band). The spacing between the tones (Df) is simply 8 MHz divided by 2560 that translates to 3.125 KHz. The tones may be further organized into Tone Sets each with four tones, and Tone Partitions, each with 20 Tone Sets. Alternatively the tones may be organized into Tone Clusters each with 20 tones, and Traffic Partitions, each with 4 Tone Clusters. A traffic channel requires at least one traffic partition. Control and access channels may be interspersed among the traffic channels in the 5 MHz slots. As will be discussed further below, data is redundant over a tone set.

The organization of the tones also permits standardization of tone assignments to users so as to permit the contemplated calculations in an orderly fashion. For example, each user may be assigned only multiples of traffic partitions. The division of the total transmission band into sub-bands also allows for lower sampling rates and less intensive DSP requirements (since the processed band is spread over a significantly smaller bandwidth). In addition, the partitions provide a convenient division for reducing the dimensionality of received vectors. This could be accomplished by combining selected tone set values (i.e., the corresponding tone set values in each cluster set). Although this involves a reduction in the number of degrees of freedom, such a tradeoff can be advantageous in systems wherein the maximum number of degrees of freedom are not necessary to accurately decode the data. Thus, by reducing the dimensionality of the tone set vector, the processing cost is significantly reduced.

As indicated above, to ensure that the signals modulated onto the separate tones do not mutually interfere by overlapping with other tones the tones set are spaced at intervals of 1/T, the symbol rate. Of course, some distortion occurs during transmission so that some interference may occur which may be removed with additional error correcting techniques.

The Use of DMT-SC

As noted above, the signals may be initially spread over assigned tones using appropriate codes or weights. These codes may be orthogonal within a given spatial cell, and may be

randomly assigned to the same tone bins within adjacent spatial cells. Thus, spreading codes may be reused in adjacent spatial cells and also may have a random correlation between adjacent spatial cells. Although the initial code assignments made by the base station may be orthogonal, it will be understood that in response to the weight adjustments made during adaptive equalization, the spreading codes will typically evolve, or adapt, to non-orthogonal codes after the communications network has been active for some time. As will be discussed in greater detail below, the criterion for the spreading codes used within a given spatial cell is advantageously linear independence rather than orthogonality. The random correlation of spreading codes in adjacent spatial cells is compensated for by means of an automatically implemented code nulling technique that nulls out correlated portions of the transmitted signals using linear weighting.

Once the spreading codes associated with the DMT-SC modulation technique have been assigned to the encoded data signals, as represented by the block 1013 of FIGURE 8, the processed signals are linearly summed, as indicated by a summing block 1025. A similar signal processing procedure is used on other incoming signals as indicated by the blocks 1021-1023, that correspond to the blocks 1011-1013. These signals are summed within the summer 1025 and assigned to carrier frequencies within the 5 MHz sub-bands shown in FIGURE 7, as indicated by a block 1030.

As noted above, during the course of adaptive equalization, the codes typically become non-orthogonal in order to maximize the SINR throughout the overall communications network 100. However, in order to retain the maximum number of degrees of freedom throughout the communications system 100, it is preferable to maintain linear independence of the complex spreading vectors throughout a given spatial cell. Linearly independent complex vectors are those that cannot be expressed as a sum or scalar multiple of any combination of the other complex vectors in the system. Thus, by preserving linear independence among the spreading codes, a matrix set of linear equations can be derived that allows each of the system variables (i.e., data symbols) to be uniquely decoded. Insofar as the spreading codes become more linearly dependent, the ability to discriminate amongst data symbols becomes more difficult. However, in some applications, band pass filter values are established in the beginning and thereafter the system must operate within those constraints.

The spread signals are linearly added on a carrier-by-carrier basis to obtain the overall DMT-SC waveform. In order to despread this signal, the received signal is detected and converted into matrix form. The received vector is multiplied by a scaling factor (that is proportional to the number of bits in the spreading code), and a matrix comprised of the spreading codes. The resultant vector provides the despread data symbols as an output. From this example, it can be induced that as many data bits can be distinctly despread as there are bits in the spreading codes, so long as the spreading codes remain linearly independent

Returning to the spreading of the data, once the encoded, spread-spectrum signals have been assigned to the frequency carrier bands, the signal may be transformed from the discrete frequency domain to the analog time domain using an inverse fast Fourier transform (IFFT) and an analog to digital converter. By using an IFFT and an FFT to provide for OFDM, multiple modulators are not required, as is well known in the art. This is because the calculations relating to the DMT-SC modulation technique are less intensive in the frequency domain than in the time domain. For this reason, the bulk of the signal processing is preferably performed in the frequency domain (with the exception of the modem operations of, for example, encryption, filtering, etc.) and is transformed to the time domain as one of the last steps before transmission.

Depending upon bandwidth considerations, signals from the same user are assigned one or more spreading codes and one or more traffic partitions. The assignment of different spreading codes and additional traffic partitions to provide additional bandwidth for a requesting user unit (a unit that communicates via one of the remotes) is particularly elegant (in comparison with bandwidth allocation using TDMA) from an implementation standpoint. This is because the allocation of new spreading codes and tone sets is mathematically simple and merely requires a numerical change to the despreading vector (for the reassignment of a new tone set) or an increase or decrease of the bandwidth of a bandpass filter on the receiving side (for the reassignment of a new tone set).

Advantages Associated with DMT-SC

The use of DMT-SC is highly advantageous in the system of the present invention. For example, the use of DMT-SC allows the channel characteristics to be evaluated at discrete points that can be exactly represented in matrix form as a complex vector. Thus, because selected tones within each tone set can be designated as pilots distributed throughout the frequency band, a simple evaluation of a finite number of complex values results in an accurate channel estimation. Furthermore, theoretically, the channel distortion can be compensated at the discrete tone frequencies by a simple complex conjugate multiplication. That is, since discrete tones are used, it is not necessary to know the entire channel response between the tones since the channel only affects operations at the exact points of the tone set frequencies. If the channel is defined at these discrete points, the received tones need only be multiplied by the appropriate complex, amplitude and phase to equalize the channel. This means that exact equalization is accomplished by a simple complex multiplication. This channel equalization calculation may be subsumed in the calculation of despread/spread weights that improve or optimize characteristics of the signal such as the signal to noise and interference ratio.

Also, the use of DMT-SC ensures that the equalization of antenna array time dispersion is very simple. In multiple element antenna arrays, a time delay is observed between receptions of a

waveform by the spatially separated sensors when the wave impinges on the array. In a very wide band system, this delay creates dispersion. However, by using DMT-SC, the dispersion can be represented by discrete values of a scaleable vector since the response is only evaluated at discrete points of the frequency.

Furthermore, each user on the system could operate with a different QAM (or other M-ary) constellation size. This is because the symbols are not spread over the entire bandwidth as in direct sequence spread spectrum. Rather, in DMT-SC the symbols are spread over frequency bins of various sizes so that each user can have the optimum size QAM constellation (i.e., the highest order allowable in a given SINR). This increases the overall system capacity since the system is not restricted to the lowest common denominator (i.e., the QAM or M-ary constellation size at which all channels can operate). In addition, at lower constellation sizes a lower signal-to-noise ratio is required to demodulate the signal, and this lower signal-to-noise ratio requirement can be used to extend the range of the base station that provides additional system flexibility.

The use of DMT-SC modulation also provides several unexpected advantages when used in combination with certain communication technologies. First of all, since DMT-SC spreading allows for flexible spreading bandwidths and gain factors (i.e., a given signal can be spread over as much bandwidth as desired), it is particularly advantageous for exploiting the spectral diversity of the channel. That is, since the channel has certain bands with better response than other bands, signals can be selectively spread over the more desirable bands.

In addition, DMT-SC also allows for the use of code-nulling to greatly improve the reuse capacity of the communication link beyond the reuse capacity of conventional CDMA. Since DMT-SC is used instead of direct sequence or frequency hopping, selected portions of the spreading code can be nulled within the despreader. Thus, only those portions of the spreading code which are not common with the interfering spreading codes will be despread. Furthermore, DMT-SC is particularly advantageous when implemented within a variable bandwidth system since the allocation of bandwidth is highly flexible in such a system, and can be implemented by the appropriate assignment of additional tones to the requesting user. In summary, DMT-SC provides a solution that nulls the interfering signals.

Finally, DMT-SC is advantageous as applied to a multi-element antenna array system where matrix calculations comprise the bulk of the processing operations. As is well known in the art, as the dimensionality of a matrix grows, the calculation operations necessary to invert the covariant matrix increases as the cube of the matrix dimensionality. Thus, the processing power increases as the cube of the matrix dimensionality and, consequently, so does the cost of the processing circuitry. Thus, in order to avoid skyrocketing costs, it is advantageous to limit the dimensionality of the matrices used to perform the spreading and despreading calculations. Since

in a multi-element antenna array system it is sometimes desirable to change the number of antenna sensor elements to enhance the beam forming capability of the system, such a system would normally incur an increase in matrix dimensionality (since each sensor corresponds to an element in the matrix). However, in a DMT-SC system, if sensors are added to the antenna array, the dimensionality of the matrix can be preserved by reducing the number of tones in each tone set.

This preservation of matrix dimensionality is possible because the mathematical formalism used when performing amplitude and phase weighting of the signal on each of the sensors is substantially similar to the formalism used when performing amplitude and phase weighting of each of the tones in a tone set. Thus, an analogy exists between the multiple sensors in an antenna array and the multiple tones in a tone set. Consequently, the same matrix can be used to determine weights for both sensor elements and tones, so that if the number of sensor elements increases, the number of tones can be decreased to compensate (i.e., preserve the same matrix dimensionality), and vice versa. Furthermore, essentially the same SINR is preserved in such a system since the degrees of freedom lost in the number of tones is regained in the number of beams. In contrast, direct sequence spread spectrum could not change the number of tones as beams are added since there are no tones to add or subtract. Thus, the cost of such a system would increase enormously relative to the cost of the system of the present invention as capacity is increased. Specifically, the cost of the present invention increases approximately proportionally with the capacity, while the cost of another system using, for example, direct sequence spread spectrum, increases as the cube of the capacity.

Once the signal has been DMT-SC modulated, the signal is output to the antenna for transmission. DMT-SC enables the appropriate signals to be directed to the appropriate user units (i.e., by means of antenna beam-forming discussed below).

Beam Forming

In accordance with one aspect of the present invention, adaptive antenna arrays are used in conjunction with a beam forming algorithm to achieve spatial diversity within each spatial cell and implement SDMA. That is, signals output by the antennas are directionally formed by selectively energizing different antenna sensors with different signal gains so that remote terminals in one portion of a spatial cell are able to communicate with the base station while other remote terminals in a different portion of the spatial cell may communicate with the same base station, even if they are using the same tone set and code. It should be understood that in the fixed implementation of the current invention, i.e., where the remote access terminals do not move substantially during communication with the base station, usually staying within a spatial cell during communication, the beam forming algorithm used in the airlink need not account for mobile remote units leaving and entering the spatial cell. In one advantageous embodiment, each

spatial cell is partitioned into four sectors where each sector transmits and receives over one of the four sub-band pairs.

As set forth above, the beam forming method of the present invention, like the use of codes, should not be conceived as separate from the overall adaptive equalization method of the present invention. Rather, the method used to selectively energize the antenna sensors (during transmission) or selectively weight the signals received on the different sensor elements (during reception) is subsumed into the overall method used to maximize SINR. The relation of the beam forming method to the overall maximization of SINR method will be described in greater detail below.

Code-Nulling

The use of spread-spectrum technology (particularly DMT-SC) and directional antennas within the preferred airlink of the present invention allows for several error cancellation benefits, including effects that are analogous to code nulling and null steering, by means of linear weighting in code and space.

Code-nulling is used to discriminate between non-orthogonal signals emanating from adjacent spatial cells. Again, the code-nulling method should be understood in the context of the maximization of SINR method of the present invention. That is, the code-nulling method should be considered as the portion of the method that maximizes SINR with respect to the code domain. This way of understanding the code-nulling method will be described in further detail with respect to FIGURE 10.

It should be understood that if signals generated within the same spatial cell or beam all have orthogonal spreading codes, code-nulling is typically not necessary since the orthogonality is sufficient to ensure that there is no cross modulation. However, as mentioned above, the spreading codes used within a particular spatial cell may not be orthogonal, although they are preferably linearly independent. Furthermore, the transceivers within the neighboring spatial cells may employ spreading codes that have a random correlation with the spreading codes used in the local spatial cell.

By adjusting the spreading weights associated with each communications channel the base station is able to cross-correlate these signals on the same tone set to subtract out interference due to "neighboring" signals. In one embodiment, the base station has the spreading codes used to spread different signals assigned to the same tone set, so that this information can be used to initially calculate the appropriate weights for nulling out interference from other codes.

As discussed above, when the spreading codes used to spread distinct data signals are orthogonal, the spread data can be precisely recovered during despreading. However, when the spreading codes are not orthogonal (as is the case with spreading codes that are used in

neighboring spatial cells), cross modulation may result so that the data signals are not able to be precisely distinguished by simple despreading (i.e., despreading without code-nulling).

In order to compensate for this phenomenon, code-nulling weights are used in the despreader. By nulling out the cross modulation present in the received signal, the appropriate values of the data bits are output by the receiver. As long as the complex spreading weights are linearly independent, and the SNR is sufficiently high, the exact symbol values can be discriminated by this method. It will be appreciated that the code-nulling procedure above is inherently implemented during derivation of the overall weights that maximize the SINR.

Null-Steering

In addition to code-nulling, an exemplary directional antenna shown in FIGURES 11 and 12 with no spectral spreading, forms signals including null regions (i.e., regions where the antenna attenuates incoming signals or where there is a very low antenna gain). These null regions can be formed in a pattern so that the nulls are directed towards known interferers (e.g., interfering signal sources or interfering multi-path reflectors). In this manner, interfering signals are de-emphasized in the spatial domain. As will be discussed in greater detail below, the use of null-steering in conjunction with code-nulling is highly advantageous.

In accordance with one aspect of the present invention, significant processing time and sophistication can be saved since significant similarity exists between the methods for performing null-steering and code-nulling. Specifically, the mathematical formalism used to achieve null-steering is analogous to the formalism used to achieve code-nulling. According to this analogy, just as the tones in a tone set are multiplied by complex weights to alter the amplitude and phase of the tones, so are the gain and relative phase of signals output and received by the antenna elements altered by a set of multiplicative weights. This multiplication by complex weights can be expressed in a matrix form for both code nulling - a spectral concept - and null steering - a spatial concept. Thus, the calculations performed in the spectral code domain correspond formally to the calculations performed in the spatial domain. Consequently, null steering can be performed in a system using code-nulling simply by adding an extra dimension to the matrices used for calculating the complex weights and multiplying the signals by these weights.

FIGURE 10 generally depicts how weights calculated in both the code and spatial domain are used to maximize the SINR. It should be noted that FIGURE 10 is primarily a conceptual representation and is not meant to convey the actual processing steps that occur in the method of maximizing SINR. As shown in FIGURE 10, a three dimensional graph plots the relationship among code, space, and SINR. Specifically, the code and spatial domains are shown in one plane, while the SINR is plotted perpendicular to the plane defined by the code and spatial domains.

The SINR is plotted on a scale of 0 to 1 where a value of 0 indicates that the signal consists entirely of noise and interference while a value of 1 indicates that the signal consists entirely of the signal of interest.

The code domain axis of the graph represents the various weighting values that can be applied to each of the tones, while the spatial domain axis of the graph represents the weighting values that can be applied to each of the antenna elements. As can be seen from the plot of FIGURE 10, certain weights applied in the correct combination of code and spatial values result in SINR values near 1 so that optimal signal detection is achieved by calculating the code and spatial weights that converge to the "peaks" depicted in FIGURE 10. The method of altering the code and spatial domain weights so that convergence to the peak SINR is achieved is described in greater detail below with reference to the method of maximizing SINR section. The invention combines spatial and spectral spreading and despreading to optimally remove interference from the received signals.

Returning to the null steering procedure that forms a portion of the method for calculating weights in the spatial domain, the null steering method, illustrated schematically in FIGURE 13, provides for increased user capacity for each base station. As depicted in FIGURE 13, a first beam, "beam A," is directed by the antenna 120 using beam-forming techniques, over a particular spatial region (i.e., the signal strength is significant in the depicted region enclosed by solid lines). A second beam, "beam B," is directed by the antenna 120 over a different spatial region (enclosed by the dashed line in FIGURE 13). Both signals include sidebands, that normally would generate interference within the adjacent signal space, and null regions between the main beam and the sidebands. Of course, it will be appreciated that more complicated beam patterns may be employed having several sidebands and null regions.

In accordance with one embodiment of the invention, the null regions of beams A and B are positioned in the direction of each of the interfering transceivers (e.g., transceivers operating on the same tone set and/or code as the intended transceiver). Thus, as depicted in FIGURE 13, while beam A is directed towards remote A (since remote A is the intended receiver) the null of beam A is directed towards remote B (since remote B is an interferer). Similarly, beam B is directed towards remote B (since remote B is the intended receiver) while the null of beam B is directed towards remote A (since remote A is an interferer). A similar weighting scheme is observed when the remotes are transmitting and the base station is receiving. The same null-steering principle also may be applied to reduce the interference due to neighboring base stations.

It should be noted here that multi-path reflectors may also be treated as interfering signal sources so that null regions can be positioned to null out signals from these reflectors. However, in one embodiment, if the reflectors are not significantly time varying, the reflected interferers are

not nulled. Rather the reflected signals are advantageously phase shifted to provide constructive interference so that the SINR is increased.

The null resolution (i.e., the closeness in degrees of the nulls) which the antenna arrays are capable of providing is dependent upon several factors. Two main factors are the spacing of the antenna sensor elements and the S/N ratio of the incoming signal. For instance, if the aperture size is sufficiently large (e.g., if the sensor elements are sufficiently far apart) then a better null resolution will result. Also, if the S/N ratio of the received signal of interest is high enough, then the signal of interest could actually be placed partially within a null (so that some gain of the signal is lost, but the overall ratio between the gain null on the interferer and the gain null on the signal of interest allows for effective cancellation of the interferer and detection of the signal of interest). For example, if 15 dB of gain is necessary to close the link for a given channel, and the S/N ratio of the signal of interest is 30 dB, while the S/N ratio of the interferer is 60 dB, then if a null of -70 dB is placed on the interferer, while the signal of interest is in the same null at about -15 dB, then the interferer will have a net -10 dB gain and the signal of interest will have a net 15 dB gain so that the interferer is canceled and the link is closed. Thus, a higher S/N ratio allows the nulls to be placed closer to the signals of interest so that a higher null resolution is achieved. It should be noted here that, in accordance with one advantageous embodiment of the invention, the depth of a given null is proportional to the strength of the interferer that is to be canceled. In addition, due to the frequency diversity provided by the system, nulls can be positioned relatively close to each other if the steering vectors (associated with the code weights) of two interfering remotes are sufficiently distinct to provide the necessary processing gain to close the communications link.

In an alternate embodiment, the remote terminals also include directional antennas in one preferred embodiment so that the remote terminals are also capable of null steering. FIGURE 16 is a graph plotting antenna gain (measured in decibels) versus direction (measured in degrees). A number of base stations are represented in FIGURE 16 by crosses, while other remotes (having non-orthogonal codes) are represented by small circles.

In the worst case scenario, the remote is located equidistant from three base stations (i.e., on a vertex of a hexagonal spatial cell). This case is represented in FIGURE 16 by the presence of three crosses that transmit with substantially equivalent signal strength. These base stations are shown at approximately 0, 90°, and -90° from the zero direction of the remote antenna.

Normally, each of the base stations would be received at the same level (i.e., at -85 dB) so that substantial interference would result between the three base stations when received at the remote. However, due to the beam forming weights applied by the directional antenna of the remote, the interfering base stations (i.e., the stations at ±90°) are attenuated by approximately 50 dB (i.e., 120 dB minus 70 dB) relative to the intended base station (i.e., the base station at 0°).

Thus, due to the fact that the beam from the receiving remote antenna is formed to have maximum gain at the intended base station, and to have minimum gain (nulls) at the strongest interfering base stations, the remote terminals are able to more easily discriminate between the signal of interest and interfering signals. That is, by means of beam forming and null-steering employed at the remote terminals a much higher signal-to-interference plus noise ratio (SINR) can be obtained in much the same manner as with the base stations.

It should be noted here that the remote terminals may also employ code nulling. In an alternate embodiment, initial code nulling weights are calculated within the base station and transmitted to the remote terminals. The remote terminals subsequently adapt the transmitted weights to maximize the SINR as required by the particular interference environment of each remote. By calculating the initial weights and sending these to the remote terminals, much of the intensive calculations need not be performed within the remotes. Thus, the remote terminals can be made more cost effectively.

In one aspect of the invention - referred to as "retrodirectivity" - the base stations adapt the spreading and despreading weights used within the base stations for transmitting and receiving signals in order to maximize the overall SINR within the communications network 100. In an alternate embodiment, this may be performed, for example, by monitoring the average bit error rate (BER) throughout the communication network 100 and modifying the spreading weights at each of the base stations, as well as each of the remote terminals , to decrease the BER.

Despread Weight Adaptation Algorithm

In one embodiment of the present invention, during the traffic establishment phase, a series of pilot tones having known amplitudes and phases, are transmitted over the entire frequency spectrum. The pilot tones are at a known level (e.g., 0 dB), and are spaced apart by approximately 30 KHz to provide an accurate representation of the channel response (i.e., the amplitude and phase distortion introduced by the communication channel characteristics) over the entire transmission band. To compensate for the channel distortion, a complex inverse (having an amplitude component and a phase component) of the channel response is calculated and multiplied by the incoming signals. This initializes the weights during the traffic establishment phase.

In certain cases, where the channel induced fade is too deep to provide an adequate signal-to-noise ratio, the tone clusters where these deep nulls occur are excised (i.e., discarded so as to not factor into the signal during despreading).

Since the channel response varies over time, the set of complex conjugate compensation weights are periodically recalculated to insure an accurate channel estimation.

Another method of channel equalization involves equalizing the channel effects (due, for example, to noise and known interferers) by data directed methods. That is, rather than transmitting a known training signal (such as a set of pilot tones), weights are applied to the received signal so as to detect a selected property of the data signal. For example, if a PSK modulation technique is used on the data, a constant power modulus is expected in the received signal. Alternately, in a QAM signal, the data will be detected in an amplitude-phase signal constellation plane to have substantially concentric rings. Thus, if the channel is equalized in such a manner as to obtain the desired signal characteristics, there is a high probability that the transmitted symbols will be accurately decoded at the receiver. This general techniques is referred to as a property restoral technique. In one embodiment of the invention the property that is restored is the finite alphabet of the QAM or M-PSK symbol.

Of course, it will be appreciated by those skilled in the art that although the channel equalization method used in accordance with the invention is conceptually separable from other signal weighting and decoding methods of the present invention (discussed below), the channel equalization method may implicitly include multiple cancellation and despreading methods. Therefore, the adaptive channel equalization method of the present invention used to maximize the SINR should not be considered as a separate method from the additional methods described below that refer to interference cancellation and signal despreading and decoding methods. Rather, the adaptive channel equalization method of the present invention should be understood to encompass a plurality of the below described methods.

Reciprocity and Retrodirectivity

TDD is particularly advantageous in the practice of this invention since with the use of TDD the linear weighting coefficients used to compensate for channel interference during transmission and reception of the encoded signals need not be re-calculated within a station. The short time duration between transmission and reception by the base station, the fact that the transmission and reception occurs in the same frequency band and only slightly separated in time (TDD), and the fact that the remote access terminals are stationary with respect to the base stations assures that the channel is approximately reciprocal. That is, the properties of the air channel between the base and the remote terminals (i.e., those properties that introduce distortion in the transmitted signal) are substantially the same for both reception and transmission. Thus, substantially the same weights can be used at a station for both despreading a signal at reception and for spreading a signal at transmission. In accordance with this retrodirectivity principle, the base station can perform most of the computation for transmission spreading weights when it computes the despreading weights on reception. The transmission spreading weights are merely scalar multiples of the reception despreading weights. Similarly, in accordance with this

retrodirectivity principle, the remote station can perform most of the computation for its transmission spreading weights when it computes its despreading weights on reception.

In an alternate embodiment of the invention, the base station can transmit the weights to the remote stations to be used in the next reception at the remote station. In this manner, processing is reduced within the remote stations since a large portion of the intensive calculations are performed solely within the base station. Thus, instead of being prohibitively sophisticated, the remote terminals can be made at a suitable size and at a reasonable expense.

Because each remote terminal stands in a different spatial relation to the other remotes and bases within the communications network, each remote terminal advantageously uses equalization weights that are individually set to maximize the SINR of signals transmitted to and received from the base station to which the remote is assigned. This may be accomplished in different ways. For instance, the base station may pre-emphasize the signals sent to the remote by a calculated set of weights. Since the pre-emphasis approximately compensates for channel distortion, the remote need not perform weight adjustment calculations that are as intensive as those calculated by the base station. Thus, the remotes need not include prohibitively sophisticated processing circuitry to implement this feature of the invention.

In one aspect of the invention optimum transmit weights are calculated based on the signals received at the base station. This is called retrodirectivity. When retrodirective adaptive equalization is used to determine the set of weights used in both reception and transmission, network-wide retrodirective adaptive equalization is accomplished. Thus, the channel characteristics throughout the entire system are accounted for in accordance with this aspect of the present invention.

Of course, as with other aspects of the invention, it will be appreciated that the reciprocity and system-wide retrodirective aspects of the present invention may also have application in a mobile environment. Specifically, if the time duration between transmission and reception in the TDD system is made sufficiently small, the channel may also be reciprocal for mobile transceivers so that the same principles set forth above apply in the mobile environment.

Zone Control

In a particularly preferred embodiment of the invention, a zone controller could be used to minimize the risk of interference between remote terminals that are near one another in adjacent spatial cells. According to this aspect of the invention, the zone controller is informed of the locations of each of the remotes and base stations within an assigned zone. Those remote terminals that are likely to interfere are assigned different codes and tone sets to minimize the risk of interference.